

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	41237	(prior or before or earlier or pre-\$6) with (since after or post-\$7 or later or subsequent) with (change or modification or alteration or update)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/28 14:58
L2	20063	(node or server or client or cluster) adj3 (state or condition)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/28 14:59
L3	31	1 same 2	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/28 15:08
L4	5541	(fault or failure or problem or error) with cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/28 15:08
L5	113	1 and 4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/28 15:13
L6	4	("5581753") or ("6557111").PN.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/11/28 15:24
L7	1	"09/899662"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/11/28 15:24
S1	1136	(714/4).CCLS.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/11/28 14:54
S2	879	management adj (cluster or node)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/11/18 17:14
S3	146797	(valid or functional or working) with (state or condition)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/18 17:16

S4	90	S2 and S3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/18 17:16
S5	2	("5392209").PN.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/11/19 11:44

**Welcome to IEEE Xplore®**

- Home
- What Can I Access?
- Log-out

**Tables of Contents**

- Journals & Magazines
- Conference Proceedings
- Standards

**Search**

- By Author
- Basic
- Advanced
- CrossRef

**Member Services**

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

**IEEE Enterprise**

- Access the IEEE Enterprise File Cabinet

[!\[\]\(aff7c69c44a5e015f18c35867ef3f5c3\_img.jpg\) Print Format](#)

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)



This Page Blank (uspto)

- Home
- What Can I Access?
- Log-out

## Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

## Search

- By Author
- Basic
- Advanced
- CrossRef

## Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

## IEEE Enterprise

- Access the IEEE Enterprise File Cabinet

## Print Format

Your search matched **114** of **1097671** documents.

A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

## Refine This Search:

You may refine your search by editing the current search expression or entering a new one in the text box.

Check to search within this result set

## Results Key:

**JNL** = Journal or Magazine **CNF** = Conference **STD** = Standard

**1 NCAPS: application high availability in Unix computer clusters**

*Laranjeira, L.A.;*  
Fault-Tolerant Computing, 1998. Digest of Papers. Twenty-Eighth Annual International Symposium on, 23-25 June 1998  
Pages:441 - 450

[\[Abstract\]](#) [\[PDF Full-Text \(44 KB\)\]](#) **IEEE CNF**

**2 Reliability analysis of clustered computing systems**

*Mendiratta, V.B.;*  
Software Reliability Engineering, 1998. Proceedings. The Ninth International Symposium on, 4-7 Nov. 1998  
Pages:268 - 272

[\[Abstract\]](#) [\[PDF Full-Text \(64 KB\)\]](#) **IEEE CNF**

**3 Experience with biased random pattern generation to meet the demand for a high quality BIST**

*Gruetzner, M.; Starke, C.W.;*  
European Test Conference, 1993. Proceedings of ETC 93., Third, 19-22 April 1993  
Pages:408 - 417

[\[Abstract\]](#) [\[PDF Full-Text \(568 KB\)\]](#) **IEEE CNF**

**4 Fault tolerant cluster computing through replication**

*Kam Hong Shum;*  
Parallel and Distributed Systems, 1997. Proceedings., 1997 International Conference on, 10-13 Dec. 1997  
Pages:756 - 761

[\[Abstract\]](#) [\[PDF Full-Text \(504 KB\)\]](#) **IEEE CNF**

**5 Reliability measurement of fault-tolerant onboard memory system under fault clustering**

*Choi, A.; Park, N.; Meyer, F.J.; Lombardi, F.; Piuri, V.;*  
Instrumentation and Measurement Technology Conference, 2002. IMTC/2002. Proceedings of the 19th IEEE, Volume: 2, 21-23 May 2002  
Pages:1161 - 1166 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(597 KB\)\]](#) **IEEE CNF**

**6 A backroom mission operations center for TechSat 21**

*Zetocha, P.;*  
Aerospace Conference Proceedings, 2002. IEEE, Volume: 7, 9-16 March 2002  
Pages:7-3189 - 7-3193 vol.7

[\[Abstract\]](#) [\[PDF Full-Text \(462 KB\)\]](#) **IEEE CNF**

**7 Application of fuzzy equivalent matrix for fault diagnosis of oil-immersed insulation**

*Zhang, G.; Yasuoka, K.; Ishii, S.; Li Yang; Zhang Yan;*  
Dielectric Liquids, 1999. (ICDL '99) Proceedings of the 1999 IEEE 13th  
International Conference on , 20-25 July 1999  
Pages:400 - 403

[\[Abstract\]](#) [\[PDF Full-Text \(260 KB\)\]](#) [IEEE CNF](#)

---

**8 Lazy garbage collection of recovery state for fault-tolerant  
distributed shared memory**

*Sultan, F.; Nguyen, T.D.; Iftode, L.;*  
Parallel and Distributed Systems, IEEE Transactions on , Volume: 13 , Issue: 7  
, July 2002  
Pages:673 - 686

[\[Abstract\]](#) [\[PDF Full-Text \(1287 KB\)\]](#) [IEEE JNL](#)

---

**9 Building topology-aware overlays using global soft-state**

*Zhichen Xu; Chunqiang Tang; Zheng Zhang;*  
Distributed Computing Systems, 2003. Proceedings. 23rd International  
Conference on , 19-22 May 2003  
Pages:500 - 508

[\[Abstract\]](#) [\[PDF Full-Text \(333 KB\)\]](#) [IEEE CNF](#)

---

**10 Comparison of feature extractors on DC power system faults for  
improving ANN fault diagnosis accuracy**

*Momoh, J.A.; Oliver, W.E., Jr.; Dolce, J.L.;*  
Systems, Man and Cybernetics, 1995. 'Intelligent Systems for the 21st  
Century'., IEEE International Conference on , Volume: 4 , 22-25 Oct. 1995  
Pages:3615 - 3623 vol.4

[\[Abstract\]](#) [\[PDF Full-Text \(552 KB\)\]](#) [IEEE CNF](#)

---

**11 Dynamic mutation testing in integrated regression analysis**

*Laski, J.; Szermer, W.; Luczycki, P.;*  
Software Engineering, 1993. Proceedings., 15th International Conference on  
, 17-21 May 1993  
Pages:108 - 117

[\[Abstract\]](#) [\[PDF Full-Text \(712 KB\)\]](#) [IEEE CNF](#)

---

**12 Clustering algorithms for Bayesian fault detection in linear systems**

*Davis, M.H.A.; Lasdas, S.; Salmon, D.J.;*  
Decision and Control, 1991., Proceedings of the 30th IEEE Conference on  
, 11-13 Dec. 1991  
Pages:674 - 675 vol.1

[\[Abstract\]](#) [\[PDF Full-Text \(144 KB\)\]](#) [IEEE CNF](#)

---

**13 Development of a testbed for distributed satellite command and  
control**

*Zetocha, P.; Brito, M.;*  
Aerospace Conference, 2001, IEEE Proceedings. , Volume: 2 , 10-17 March  
2001  
Pages:2/609 - 2/614 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(572 KB\)\]](#) [IEEE CNF](#)

---

**14 Detection of incipient fault using fuzzy agglomerative clustering  
algorithm**

*Boudaoud, N.; Masson, M.;*  
Fuzzy Information Processing Society, 1999. NAFIPS. 18th International  
Conference of the North American , 10-12 June 1999  
Pages:233 - 237

**15 A fault-tolerant rearrangeable permutation network**

*Yang, Y.; Wang, J.;*

Computers, IEEE Transactions on , Volume: 53 , Issue: 4 , April 2004

Pages:414 - 426



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used point to point cluster state fault

Found 25,403 of 147,060

Sort results by

relevance

[Save results to a Binder](#)[Try an Advanced Search](#)

Display results

expanded form

[Search Tips](#) [Open results in a new window](#)[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

## 1 Understanding fault-tolerant distributed systems

Flavin Cristian

February 1991 **Communications of the ACM**, Volume 34 Issue 2Full text available: [pdf\(6.17 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

## 2 Research sessions: consistency and availability: Highly available, fault-tolerant, parallel dataflows

Mehul A. Shah, Joseph M. Hellerstein, Eric Brewer

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**Full text available: [pdf\(210.17 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

We present a technique that masks failures in a cluster to provide high availability and fault-tolerance for long-running, parallelized dataflows. We can use these dataflows to implement a variety of continuous query (CQ) applications that require high-throughput, 24x7 operation. Examples include network monitoring, phone call processing, click-stream processing, and online financial analysis. Our main contribution is a scheme that carefully integrates traditional query processing techniques for ...

## 3 Fault tolerance: Collective operations in application-level fault-tolerant MPI

Greg Bronevetsky, Daniel Marques, Keshav Pingali, Paul Stodghill

June 2003 **Proceedings of the 17th annual international conference on Supercomputing**Full text available: [pdf\(206.31 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Fault-tolerance is becoming a critical issue on high-performance platforms. Checkpointing techniques make programs fault-tolerant by saving their state periodically and restoring this state after failure. *System-level* checkpointing saves the state of the entire machine on stable storage, but this usually has too much overhead. In practice, programmers do manual checkpointing by writing code to (i) save the values of key program variables at critical points in the program, and (ii) restore ...

**Keywords:** MPI, application-level checkpointing, collective communication, fault-tolerance, non-FIFO communication, scientific computing

4 Fast cluster failover using virtual memory-mapped communication

Yuanyuan Zhou, Peter M. Chen, Kai Li

May 1999 **Proceedings of the 13th international conference on Supercomputing**

Full text available: [pdf\(1.45 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Full text available: [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

6 Specifying and using a partitionable group communication service

Alan Fekete, Nancy Lynch, Alex Shvartsman

August 1997 **Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing**

Full text available: [pdf\(1.19 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 Fault-tolerance in the advanced automation system

Flaviu Cristian, Bob Dancey, Jon Dehn

September 1990 **Proceedings of the 4th workshop on ACM SIGOPS European workshop**

Full text available: [pdf\(1.39 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The Advanced Automation System is a distributed real-time system under development by IBM's Systems Integration Division for the US Federal Aviation Administration. The system is intended to replace the present en-route and terminal approach US air traffic control computer systems over the next decade. High availability of air traffic control services is an essential requirement of the system. This paper discusses the general approach to fault-tolerance adopted in AAS, by reviewing some of the q ...

8 Implementing global memory management in a workstation cluster

M. J. Feeley, W. E. Morgan, E. P. Pighin, A. R. Karlin, H. M. Levy, C. A. Thekkath

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available: [pdf\(1.52 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

9 Specifying and using a partitionable group communication service

May 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 2

Full text available: [pdf\(473.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Group communication services are becoming accepted as effective building blocks for the construction of fault-tolerant distributed applications. Many specifications for group communication services have been proposed. However, there is still no agreement about what these specifications should say, especially in cases where

the services are partitionable, i.e., where communication failures may lead to simultaneous creation of groups with disjoint memberships, such that each ...

**Keywords:** composable building blocks, conditional performance analysis, distributed algorithms, group communication protocols, message-passing protocols, ordered broadcast, service specification, total-order broadcast

## **10 Fault-tolerance in air traffic control systems**

Flaviu Cristian, Bob Dancey, Jon Dehn

August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Full text available:  [pdf\(264.57 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The distributed real-time system services developed by Lockheed Martin's Air Traffic Management group serve the infrastructure for a number of air traffic control systems. Either completed development or under development are the US Federal Aviation Administration's Display System Replacement (DSR) system, the UK Civil Aviation Authority's New Enroute Center (NERC) system, and the Republic of China's Air Traffic Control Automated System (ATCAS). These systems are intended to replace present ...

**Keywords:** exception handling, failure, failure classification, failure masking, failure semantics, fault-tolerant systems, group communications, redundancy, server group, software robustness, system architecture

## **11 The architecture and performance of security protocols in the ensemble group communication system: Using diamonds to guard the castle**

August 2001 **ACM Transactions on Information and System Security (TISSEC)**, Volume 4 Issue 3

Full text available:  [pdf\(418.73 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

Ensemble is a Group Communication System built at Cornell and the Hebrew universities. It allows processes to create *process groups* within which scalable reliable fifo-ordered multicast and point-to-point communication are supported. The system also supports other communication properties, such as causal and total multicast ordering, flow control, and the like. This article describes the security protocols and infrastructure of Ensemble. Applications using Ensemble with the extensions des ...

**Keywords:** Group communication, security

## **12 Automated application-level checkpointing of MPI programs**

Greg Bronevetsky, Daniel Marques, Keshav Pingali, Paul Stodghill

June 2003 **ACM SIGPLAN Notices , Proceedings of the ninth ACM SIGPLAN symposium on Principles and practice of parallel programming**, Volume 38 Issue 10

Full text available:  [pdf\(130.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The running times of many computational science applications, such as protein-folding using *ab initio* methods, are much longer than the mean-time-to-failure of high-performance computing platforms. To run to completion, therefore, these applications must tolerate hardware failures. In this paper, we focus on the stopping failure model in which a faulty process hangs and stops responding to the rest of the system. We argue that tolerating such faults is best done by an approach called appl ...

**Keywords:** MPI, application-level checkpointing, fault-tolerance, non-FIFO communication, scientific computing

## 13 Run-time adaptation in river

Remzi H. Arpacı-Dusseau

February 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 1

Full text available: [pdf\(849.04 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

**Keywords:** Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

## 14 Bayeux: an architecture for scalable and fault-tolerant wide-area data dissemination

Shelley Q. Zhuang, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz, John D. Kubiatowicz

January 2001 **Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video**

Full text available: [pdf\(272.26 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The demand for streaming multimedia applications is growing at an incr edible rate. In this paper, we propose Bayeux, an efficient application-level multicast system that scales to arbitrarily large receiver groups while tolerating failures in routers and network links. Bayeux also includes specific mechanisms for load-balancing across replicate root nodes and more efficient bandwidth consumption. Our simulation results indicate that Bayeux maintains these properties while keeping transmi ...

## 15 Multigrain shared memory

Donald Yeung, John Kubiatowicz, Anant Agarwal

May 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 2

Full text available: [pdf\(369.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

Parallel workstations, each comprising tens of processors based on shared memory, promise cost-effective scalable multiprocesssing. This article explores the coupling of such small- to medium-scale shared-memory multiprocessors through software over a local area network to synthesize larger shared-memory systems. We call these systems Distributed Shared-memory MultiProcessors (DSMPs). This article introduces the design of a shared-memory system that uses multiple granularities of sharing, ca ...

**Keywords:** distributed memory, symmetric multiprocessors, system of systems

## 16 Programming languages for distributed computing systems

Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum

September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

Full text available: [pdf\(6.50 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less

satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

**17 SUVS: a distributed real-time system testbed for fault-tolerant computing**

S. M. Yang, K. M. Kavi, A. Agarwalla, M. Reddy, S. Anam

March 1992 **Proceedings of the 1992 ACM/SIGAPP symposium on Applied computing: technological challenges of the 1990's**

Full text available: [pdf\(1.09 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

**18 Comparing the effectiveness of fine-grain memory caching against page migration/replication in reducing traffic in DSM clusters**

An-Chow Lai, Babak Falsafi

July 2000 **Proceedings of the twelfth annual ACM symposium on Parallel algorithms and architectures**

Full text available: [pdf\(105.85 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we compare and contrast two techniques to improve capacity/conflict miss traffic in CC-NUMA DSM clusters. Page migration/replication optimizes read-write accesses to a page used by a single processor by migrating the page to that processor and replicates all read-shared pages in the sharers' local memories. R-NUMA optimizes read-write accesses to any page by allowing a processor to cache that page in its main memory. Page migration/replication requires less hardware c ...

**19 Design challenges of virtual networks: fast, general-purpose communication**

Alan M. Mainwaring, David E. Culler

May 1999 **ACM SIGPLAN Notices , Proceedings of the seventh ACM SIGPLAN symposium on Principles and practice of parallel programming,**  
Volume 34 Issue 8

Full text available: [pdf\(1.57 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Virtual networks provide applications with the illusion of having their own dedicated, high-performance networks, although network interfaces posses limited, shared resources. We present the design of a large-scale virtual network system and examine the integration of communication programming interface, system resource management, and network interface operation. Our implementation on a cluster of 100 workstations quantifies the impact of virtualization on small message latencies and throughput ...

**Keywords:** application programming interfaces, direct network access, high-performance clusters, protocol architecture and implementation, system resource management, virtual networks

**20 Experience Using Multiprocessor Systems—A Status Report**

Anita K. Jones, Peter Schwarz

June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Full text available: [pdf\(4.48 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)